Non-positive curvature, and the planar embedding conjecture

Anastasios Sidiropoulos*

April 30, 2013

Abstract

The planar embedding conjecture asserts that any planar metric admits an embedding into L_1 with constant distortion. This is a well-known open problem with important algorithmic implications, and has received a lot of attention over the past two decades. Despite significant efforts, it has been verified only for some very restricted cases, while the general problem remains elusive.

In this paper we make progress towards resolving this conjecture. We show that every planar metric of non-positive curvature admits a constant-distortion embedding into L_1. This confirms the planar embedding conjecture for the case of non-positively curved metrics.

*Department of Computer Science, University of Illinois at Urbana-Champaign; http://www.sidiropoulos.org/; sidiropo@illinois.edu. Supported in part by David and Lucille Packard Fellowship and by NSF grants CCF-0915984 and CCF-0915519.

1 Introduction

If $(X, d_X), (Y, d_Y)$ are metric spaces, and $f : X \to Y$ is injective, the *distortion* of f is defined to be $\mathsf{distortion}(f) = \|f\|_{\mathsf{Lip}} \cdot \|f^{-1}\|_{\mathsf{Lip}}$, where $\|f\|_{\mathsf{Lip}} = \sup_{x \neq y \in X} \frac{d_Y(f(x),f(y))}{d_X(x,y)}$. For any metric space (X, d), we use $c_1(X, d)$ to denote the L_1 *distortion of* (X, d), i.e. the infimum over all numbers D such that X admits an embedding into L_1 with distortion D. For a graph $G = (V, E)$ we write $c_1(G) = \sup c_1(V, d)$ where d ranges over all shortest-path metrics supported on G, and for a family $\mathcal{F}$ of graphs, we write $c_1(\mathcal{F}) = \sup_{G \in \mathcal{F}} c_1(G)$. Thus for a family $\mathcal{F}$ of finite graphs, $c_1(\mathcal{F}) \leq D$ if and only if every geometry supported on a graph in $\mathcal{F}$ embeds into L_1 with distortion at most D.

In the seminal works of Linial-London-Rabinovich [LLR95], and later Aumann-Rabani [AR98] and Gupta-Newman-Rabinovich-Sinclair [GNRS04], the geometry of graphs is related to the classical study of the relationship between flows and cuts.

A *multi-commodity flow instance in* G is specified by a pair of non-negative mappings $\mathsf{cap} : E \to \mathbb{R}$ and $\mathsf{dem} : V \times V \to \mathbb{R}$. We write $\mathsf{maxflow}(G; \mathsf{cap}, \mathsf{dem})$ for the value of the *maximum concurrent flow* in this instance, which is the maximal value ε such that $\varepsilon \cdot \mathsf{dem}(u, v)$ can be simultaneously routed between every pair $u, v \in V$ while not violating the given edge capacities.

A natural upper bound on $\mathsf{maxflow}(G; \mathsf{cap}, \mathsf{dem})$ is given by the *sparsity* of any cut $S \subseteq V$:

$$\frac{\sum_{uv \in E} \mathsf{cap}(u,v)|\mathbf{1}_S(u) - \mathbf{1}_S(v)|}{\sum_{u,v \in V} \mathsf{dem}(u,v)|\mathbf{1}_S(u) - \mathbf{1}_S(v)|}, \tag{1}$$

where $\mathbf{1}_S : V \to \{0, 1\}$ is the indicator function for membership in S. We write $\mathsf{gap}(G)$ for the maximum gap between the value of the flow and the upper bounds given by (1), over all multi-commodity flow instances on G. This is the *multi-commodity max-flow/min-cut gap for* G. The fundamental connection between embeddings into L_1 and multi-commodity flows is captured in the following result.

Theorem 1.1 ([LLR95, GNRS04])**.** *For every graph* G*,* $c_1(G) = \mathsf{gap}(G)$.

In particular, combined with the techniques of [LR99, LLR95], this implies that for any graph G, there exists a $c_1(G)$-approximation for the general Sparsest Cut problem.

1.1 The planar embedding conjecture

It has been shown by [LLR95, LR99] that for general graphs, $c_1(G) = \Omega(\log n)$, and there has since been a lot of effort in trying to prove that $c_1(G)$ is bounded by some universal constant for interesting classes of graphs. The most well-known open case is the so-called *planar embedding conjecture*, summarized in the following:

Conjecture 1 (Planar embedding conjecture)**.** *For every planar graph* G*,* $c_1(G) = O(1)$.

Despite several attempts on resolving this question, there has only been very little progress. More specifically, the work of Okamura & Seymour [OS81] implies that the metric induced on a *single face* of a planar graph embeds with constant distortion into L_1. In [GNRS04] it is shown that $c_1(G) = O(1)$ for any series-parallel, or outerplanar graph G. This result was extended to $O(1)$-outerplanar graphs in [CGN$^+$06]. Chakrabarti *et al.* [CJLV08] obtained constant distortion embeddings of graphs that exclude a $(K_5 \setminus e)$-minor. Note that even the case of planar graphs of treewidth 3 remains open. We remark that the best-known upper bound on $c_1(G)$ for planar graphs is $O(\sqrt{\log n})$, due to Rao [Rao99], while the best-known lower bound is 2, due to Lee & Raghavendra [LR10].

1.2 Generalizations: The GNRS conjecture

Gupta, Newman, Rabinovich, and Sinclair [GNRS04] posed the following generalization of the planar embedding conjecture, which seeks to *characterize* the graph families $\mathcal{F}$ such that $c_1(\mathcal{F}) = O(1)$, which by Theorem 1.1 also characterizes all graphs with multi-commodity gap bounded by some universal constant:

Conjecture 2 (GNRS conjecture [GNRS04])**.** *For every family of finite graphs $\mathcal{F}$, one has $c_1(\mathcal{F}) = O(1)$ if and only if $\mathcal{F}$ forbids some minor.*

We note that a strengthening of the GNRS conjecture for *integral* multi-commodity flows has also been considered [CSW13]. This is a seemingly harder problem, and progress has been even more limited in this case.

At first glance, it might appear that the GNRS conjecture is a vast generalization of the planar embedding conjecture, since planar graphs exclude K_5 as a minor. Despite this, Lee & Sidiropoulos [LS09] have shown that the GNRS conjecture is *equivalent* to the conjunction of the planar embedding conjecture, with the manifestly simpler *k-sum embedding conjecture* summarized bellow. For a graph family $\mathcal{F}$, let $\oplus_k \mathcal{F}$ denote the closure of $\mathcal{F}$ under k-clique sums (see [LS09] for a more detailed exposition). We note that the case $k = 1$ is folklore, while recently progress has been reported for the case $k = 2$ by Lee and Poore [LP13]; even for $k = 2$ however, the problem remains open.

Conjecture 3 (k-sum conjecture [LS09])**.** *For any family of graphs $\mathcal{F}$, we have $c_1(\mathcal{F}) = O(1)$ if and only if $c_1(\oplus_k \mathcal{F}) = O(1)$ for every $k \in \mathbb{N}$.*

It is therefore apparent that the planar embedding conjecture is a major step towards determining the multi-commodity gap in *arbitrary* graphs.

1.3 Our results

All previous attempts on the planar embedding conjecture have been *topological* in nature, meaning that they seek to obtain constant-distortion embeddings by restricting the topology of the planar graph. As a consequence, all known methods are insufficient even for planar graphs of treewidth 3.

We depart from this paradigm by instead restricting the *geometry* of the planar metric. For any metric (X, d), we have that (X, d) is the shortest-path metric of a planar graph if and only if it can be realized as a set of points in a simply-connected (i.e. planar) surface. We say that a planar metric is *non-positively curved* if it can be realized as a set of points in a surface of non-positive curvature (see Section 1.5 for the definition of non-positively curved spaces). This leads to a natural, and very rich class of planar metrics. For instance, non-positively curved planar metrics include all trees, all regular grids (up to constant distortion), and arbitrary subsets of the hyperbolic plane $\mathbb{H}^2$.

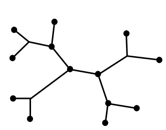
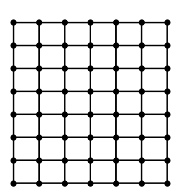
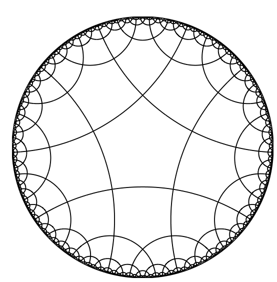

Our main result is as follows.

Theorem 1.2 (Main)**.** *There exists a universal constant $\gamma > 1$, such that every non-positively curved planar metric admits an embedding into L_1 with distortion at most γ.*

Since we are motivated by the applications of metric embeddings in computer science, we will restrict our discussion to finite metrics. We remark however that our result can be extended to obtain constant-distortion embeddings of arbitrary simply-connected surfaces of non-positive curvature into L_1.[1]

We note that embeddings of various hyperbolic spaces have been previously considered. We refer to [KL06, BS, BS05b, BS05a]. However, none of the previous results captures L_1 embeddings of arbitrary non-positively curved planar metrics. In fact, our approach is significantly different than all previous works.

1.4 A high-level overview of our approach

We now give an informal, and somewhat imprecise overview of some of the main challenges that we face when trying to embed non-positively curved planar metrics into L_1.

Distributions over monotone cuts. Let (X, d) be a metric space. We will use the standard representation of L_1 as the cone of cut pseudo-metrics (see Section 1.5 for the definition). This means that in order to embed a space into L_1 with constant distortion, it suffices to find a probability distribution over cuts, such that the probability that any pair of points x, y gets separated, is $\Theta(\alpha \cdot d(x, y))$, for some normalization factor $\alpha > 0$.

It follows by the work of Lee and Raghavendra [LR10] (see also [CJLV08]) that when seeking a constant-distortion embedding of certain spaces into L_1 it suffices to consider distributions over a specific type of cuts, called *monotone*. More precisely, let x be a fixed point. We say that a cut S is monotone (w.r.t. x) if every shortest path starting from x crosses S at most once. Let us say that a metric space is a *bundle* if there exist two points s, t, such that for every point z, there exists an s-t geodesic containing z. Then it is shown in [LR10] that a bundle admits a constant-distortion embedding into L_1 if and only if it is a convex combination of monotone cuts (i.e. a convex combination of cut pseudo-metrics, where every indicator set is a monotone cut). It is easy to show that every finite non-positively curved metric admits an isometric embedding into a bundle. We can therefore focus our efforts into finding a good distribution over monotone cuts.

The structure of monotone cuts in non-positively curved spaces. It is convenient to demonstrate the main ideas using the following example of a "pinched square". Let $X = [0, 1]^2$, endowed with the Euclidean distance. The space X can be embedded isometrically into L_1 by taking an appropriate distribution over random half-plane cuts (e.g. by choosing a uniformly random point $p \in X$, and taking the half-plane supported by a line passing through p forming a uniformly random angle with the x-axis). Let A be one of the sides of X, and let $Y = X/A$ be the quotient space obtained by contracting A into a single point, which we will refer to as the *basepoint*.

[1] This connection was pointed out by James R. Lee.

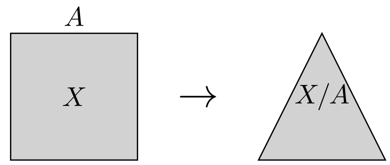

Strictly speaking, the resulting Y is not a space of non-positive curvature (in particular, there exist pairs of points in Y with two distinct geodesics joining them). However, Y admits a constant-distortion embedding into a planar surface of non-positive curvature, so in order to simplify the exposition, we may use Y without loss of generality.

It is fairly easy to see that even though Y might "look" like a triangle, its geometry is far from that of a flat Euclidean triangle. In fact, one can show that Y cannot be embedded into the Euclidean plane with bounded distortion. As a consequence, embedding Y into L_1 requires a significantly more involved distribution over cuts. Such a distribution can be constructed using cuts of the following form: For every $r \in [0,1]$, we have a family of cuts S that are contained inside the ball of radius r from the basepoint, and with boundary ∂C given by a function of period $\Theta(r)$. Roughly speaking, these cuts can be obtained by random shifts along the x-axis of cuts from the following infinite family:

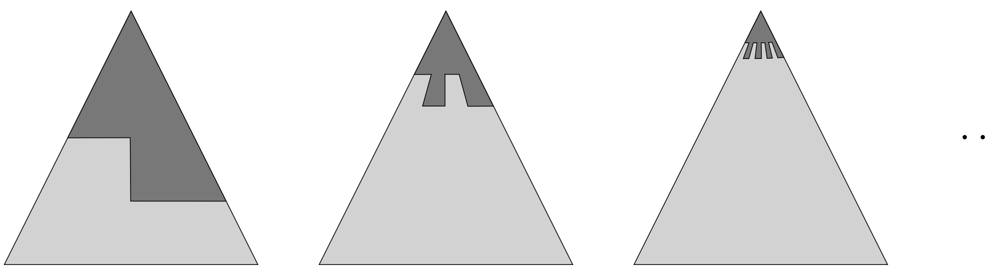

Here, the probability of a cut decreases when $r \to 0$. It is important to note that the structure of a cut depends on the distance of its boundary to the basepoint. It can be shown that this is the case for *any* constant distortion embedding of Y into L_1. Moreover, in an constant-distortion embedding, this transition has to happen in a smooth way as $r \to 0$.

Handling multiple scales. Suppose now that we modify the space Y as follows. Let R be a ray in Y, i.e. an unbounded geodesic starting at the basepoint, and let R' be a suffix of R. Cutting Y along R' introduces two copies R'_1, R'_2 of R' as segments of the boundary. We glue a copy of Y along $R'_1 \cup R'_2$ as follows:

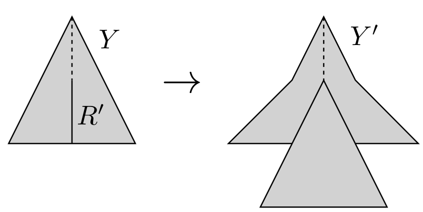

The resulting space Y' again embeds with constant distortion into a planar metric of non-positive curvature. Constructing a constant-distortion embedding for Y' requires the use of even more intricate families of cuts. Intuitively, a single cut now has to "gracefully" combine information

form multiple different scales. Let $p_1, p_2 \in Y'$ be the basepoints of the two copies of Y in Y'. The structure of a "typical" cut S has to depend on the distances between ∂S, and *both* p_1, and p_2:

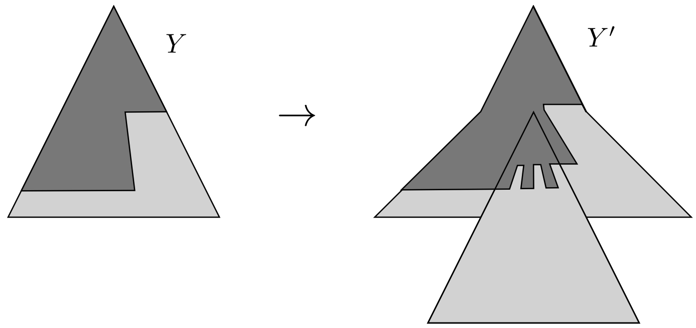

A naive way to address this problem would be to define a distribution over cuts for every single scale, and then try to combine them into a single distribution. The problem with this approach is that cuts from different scales (in our example, cuts for the two different copies of Y) might not agree on their boundary. This disagreement results in larger distortion every time we combine two different scales. Since there can be many scales, this methods leads to unbounded distortion.

We overcome this obstacle by designing a distribution over cuts that is *scale-independent*. This is done by starting with a distribution over cuts that handles large distances, and gradually modifying it to handle smaller scales. The main technical contribution of this paper is showing that in a non-positively curved surface, this can be done without increasing the distortion.

1.5 Preliminaries

We now review some basic definitions and notions which appear throughout the paper.

Graphs. Let G, and let $S \subseteq V(G)$. We denote by $G[S]$ the subgraphs of G induced by S, i.e. $G[S] = \left(S, E(G) \cap \binom{S}{2}\right)$. We will consider graphs with every edge having a non-negative length. We say that a graph is *unweighted* if all of its edges have unit length. Let $\mathsf{diam}(G)$ denote the diameter of G, i.e. $\mathsf{diam}(G) = \max_{x,y \in V(G)} d_G(x, y)$. We refer to a path between two vertices $x, y \in V(G)$ as a x-y path.

Cuts and L_1 embeddings. A cut of a graph G is a partition of $V(G)$ into $(S, \bar{S})$—we sometimes refer to a subset $S \subseteq V$ as a cut as well. A cut gives rise to a pseudometric; using indicator functions, we can write the cut pseudometric as $\rho_S(x, y) = |\mathbf{1}_S(x) - \mathbf{1}_S(y)|$. A central fact is that embeddings of finite metric spaces into L_1 are equivalent to sums of positively weighted cut metrics over that set (for a simple proof of this see [DL97]).

A *cut measure on* G is a function $\mu : 2^V \to \mathbb{R}_+$ for which $\mu(S) = \mu(\bar{S})$ for every $S \subseteq V$. Every cut measure gives rise to an embedding $f : V \to L_1$ for which

$$\|f(u) - f(v)\|_1 = \int |\mathbf{1}_S(u) - \mathbf{1}_S(v)|\, d\mu(S), \tag{2}$$

where the integral is over all cuts $(S, \bar{S})$. Conversely, to every embedding $f : V \to L_1$, we can associate a cut measure μ such that (2) holds.

Non-positively curved spaces. We will describe our proof using the definition of non-positive curvature in the sense of Busemann. We give here a brief overview of some of the relevant terminology, and we refer the reader to [Pap05, TL97] for a more detailed exposition. A metric space (X, d) is called *geodesic* if for every pair of points there exists a geodesic joining them. We say that (X, d) is non-positively curved, if for any pair of affinely parameterized geodesics $\gamma : [a, b] \to X$, $\gamma' : [a', b'] \to X$, the map $D_{\gamma,\gamma'} : [a, b] \times [a', b'] \to \mathbb{R}$ defined by

$$D_{\gamma,\gamma'}(t, t') = d(\gamma(t), \gamma'(t))$$

is convex. As we show, this property is sufficient to obtain constant-distortion embeddings of simply-connected surfaces into L_1.

Lipschitz partitions. Let (X, d) be a metric space. A distribution $\mathcal{F}$ over partitions of X is called (β, Δ)-*Lipschitz* if every partition in the support of $\mathcal{F}$ has only clusters of diameter at most Δ, and for every $x, y \in X$,

$$\Pr_{C \in \mathcal{F}}[C(x) \neq C(y)] \leq \beta \cdot \frac{d(x, y)}{\Delta}.$$

We denote by $\beta_{(X,d)}$ the infimum β such that for any $\Delta > 0$, the metric (X, d) admits a (Δ, β)-Lipschitz random partition, and we refer to $\beta_{(X,d)}$ as the *modulus of decomposability* of (X, d). The following theorem is due to Klein, Plotkin, and Rao [KPR93], and Rao [Rao99].

Theorem 1.3 ([KPR93], [Rao99]). *For any planar graph G, we have $\beta_{(V(G),d_G)} = O(1)$.*

Stochastic embeddings. A mapping $f : X \to Y$ between two metric spaces (X, d) and (Y, d') is *non-contracting* if $d'(f(x), f(y)) \geq d(x, y)$ for all $x, y \in X$. If (X, d) is any finite metric space, and $\mathcal{Y}$ is a family of finite metric spaces, we say that (X, d) *admits a stochastic D-embedding into $\mathcal{Y}$* if there exists a random metric space $(Y, d') \in \mathcal{Y}$ and a random non-contracting mapping $f : X \to Y$ such that for every $x, y \in X$,

$$\mathbb{E}\left[d'(f(x), f(y))\right] \leq D \cdot d(x, y). \tag{3}$$

The infimal D such that (3) holds is the *distortion of the stochastic embedding.* For a graph G and a graph family $\mathcal{F}$ we write $G \overset{D}{\rightsquigarrow} \mathcal{F}$ to denote the fact that G stochastically embeds into a distribution over graphs in $\mathcal{F}$, with distortion D. We also use the notation $G \rightsquigarrow \mathcal{F}$ to denote the fact that $G \overset{D}{\rightsquigarrow} \mathcal{F}$, for some universal constant $D \geq 1$. We will use the following fact.

Lemma 1.4. *Let $\mathcal{F}$ be a family of graphs, such that every $H \in \mathcal{F}$ admits an embedding into L_1 with distortion at most $\alpha \geq 1$. Let G be a graph, such that $G \overset{\beta}{\rightsquigarrow} \mathcal{F}$, for some $\beta \geq 1$. Then, G admits an embedding into L_1 with distortion at most $\alpha\beta$.*

Let G be a graph, and let $A \subseteq V(G)$. The *dilation* of A is defined to be

$$\mathsf{dil}_G(A) = \max_{u,v \in V(G)} \frac{d_{G[A]}(u, v)}{d_G(u, v)}$$

For two graphs G, G', a *1-sum* of G with G' is a graph obtained by taking two disjoint copies of G and G', and identifying a vertex $v \in V(G)$ with a vertex $v' \in V(G')$. For a graph family $\mathcal{X}$, we denote by $\oplus_1\mathcal{X}$ the closure of $\mathcal{X}$ under 1-sums.

Lemma 1.5 (Peeling lemma [LS09])**.** *Let G be a graph, and $A \subseteq V(G)$. Let $G' = (V(G), E')$ be a graph with $E' = E(G) \setminus E(G[A])$, and let $\beta = \beta_{(V, d_{G'})}$ be the corresponding modulus of decomposability. Then, there exists a graph family $\mathcal{F}$ such that $G \overset{D}{\rightsquigarrow} \mathcal{F}$, where $D = O(\beta \cdot \mathsf{dil}_G(A))$, and every graph in $\mathcal{F}$ is a 1-sum of isometric copies of the graphs $G[A]$ and $\{G[V \setminus A \cup \{a\}]\}_{a \in A}$.*

1.6 Organization

The rest of the paper is organized as follows. In Section 2 we show how to embed an arbitrary non-positively curved planar metric into an unweighted graph of special structure, called a *funnel*. In Section 3 we show how to stochastically embed a funnel into a distribution over simpler graphs, called *pyramids*. In Section 4 we introduce some of the machinery that we will use when defining our embedding into L_1. More specifically, we describe the basic operation of cuts that will allow to gradually modify a cut when computing our embedding. Using this machinery, we describe our embedding in Section 5. Finally, in Section 6 we prove that the constructed embedding has constant distortion.

2 A canonical representation of non-positively curved planar metrics

In this section we show that non-positively curved planar metrics can be embedded with constant-distortion into a certain type of unweighted planar graphs that we call *funnels*. Intuitively, a funnel is obtained by taking the union of a tree having all its leaves at the same level, with a collection of cycles, where every cycle spans all the vertices in a single layer of the tree.

Definition 2.1 (Funnel)**.** *Let G be an unweighted planar graph, and let $v \in V(G)$. We say that G is a* funnel *with* basepoint *v if the following conditions are satisfied:*

(1) There exists a collection of pairwise vertex-disjoint cycles $C_1, \ldots, C_\Delta \subset G$, such that $V(G) = \bigcup_{i=1}^{\Delta} V(C_i)$. For notational convenience, we allow a cycle C_i to consist of a single vertex, in which case it has no edges. Moreover, we have $V(C_1) = \{v\}$. We refer to each C_i as a layer *of G.*

(2) For every $i \in \{2, \ldots, \Delta - 1\}$, the graph $G \setminus V(C_i)$ has exactly two connected components, one with vertex set $\bigcup_{j=1}^{i-1} V(C_j)$, and another with vertex set $\bigcup_{j=i+1}^{\Delta} V(C_j)$.

(3) For every $i \in \{2, \ldots, \Delta\}$, every $u \in V(C_i)$ has exactly one neighbor $u' \in V(C_{i-1})$. We refer to u' as the parent *of u. In particular, v is the parent of all vertices in $V(C_2)$.*

(4) For every $i \in \{1, \ldots, \Delta - 1\}$, every $w \in V(C_i)$ has at least one neighbor $w' \in V(C_{i+1})$. We refer to every such w' as a child *of w.*

Let R be a path in G between v, and a vertex $u \in V(C_\Delta)$. We say that R is a ray*. We denote by* $\mathsf{Funnels}$ *the family of all funnel graphs. Figure 1 depicts an example of a funnel.*

We will use the following two facts about metric spaces of non-positive curvature (see e.g. [Pap05]).

Lemma 2.2. *Let $(\mathcal{S}, d)$ be a geodesic metric space of non-positive curvature. Let $x^*, x, y \in \mathcal{S}$, and let $\gamma : [0, d(x, y)] \to \mathcal{S}$ be a geodesic between x, and y. Then, the function $f : [0, 1] \to \mathbb{R}$, with $f(t) = d(x^*, \gamma(t))$ is convex.*

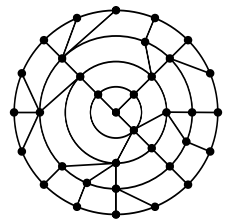

Figure 1: A funnel.

Lemma 2.3. *Let $(\mathcal{S}, d)$ be a geodesic metric space of non-positive curvature, and let $x^*, x, y \in \mathcal{S}$. Let $\gamma_x : [0, d(x^*, x)] \to \mathcal{S}$ be a geodesic between x^*, and x, and let $\gamma_y : [0, d(x^*, y)] \to \mathcal{S}$ be a geodesic between x^*, and y. Then, the function $f : [0,1] \to \mathbb{R}$, with $f(t) = d(\gamma_x(t), \gamma_y(t))$ is non-decreasing.*

Recall that for a metric space (X, d), and some $r > 0$, an *r-net* in (X, d) is a maximal subset $X' \subseteq X$ such that for any $x, y \in X'$, we have $d(x, y) \geq r$.

Lemma 2.4 (Funnel representation)**.** *Let $\mathcal{S}$ be a simply-connected surface, and let d be a non-positively curved metric on $\mathcal{S}$. Let $X \subset \mathcal{S}$ be a finite set of points. Then, (X, d) admits an embedding into a funnel with constant distortion.*

Proof. By scaling d, we may assume w.l.o.g. that the minimum distance in (X, d) is at least 1 (note that scaling d results into a metric which is still of non-positive curvature). Let $x^* \in \mathcal{S}$ be an arbitrary point. For any $x \in \mathcal{S}$, let $\gamma(x)$ denote the unique geodesic between x, and x^*. Let $r = 1/8$. For any integer $i \geq 0$, let

$$D_i = \{x \in \mathcal{S} : d(x^*, x) \leq ir\}.$$

Since d is non-positively curved, we have that for every i, the set D_i is a disk (see e.g. [Pap05]). Let Γ_i be the cycle in $\mathcal{S}$ bounding D_i. Let $\Delta = \min\{i \in \mathbb{N} : X \subset D_i\}$. Let N_Δ be an r-net in Γ_Δ. Note that since $(\mathcal{S}, d)$ is non-positively curved, there exists a unique geodesic between any pair of points. This implies that the subspace

$$T = \bigcup_{x \in N_\delta} \gamma(x)$$

is a (simplicial) tree. For every $i \in \{0, \ldots, \Delta - 1\}$, we define an r-net N_i of Γ_i as follows. Suppose that N_{i+1} is already defined. Let Y'_i be the set of points $p \in N_\Delta$ such that $\gamma(p)$ intersects N_{i+1}. Let $N'_i = \Gamma_i \cap (\bigcup_{p \in Y'_i} \gamma(p))$. Note that for any $x \in \Gamma_i$, there exists $y \in N'_i$ such that $d(x, y) < r$. Therefore, we can set to be a maximal subset $N_i \subseteq N'_i$, such that N_i is an r-net. This concludes the definition of the sequence of subsets $N_0, \ldots, N_\Delta$. Note that $N_0 = \{x^*\}$.

We define a graph G, with $V(G) = \bigcup_{i=0}^{\Delta} N_i$. The set of edges $E(G)$ is defined as follows. For every $i \in \{1, \ldots, \Delta\}$, we add a unit-length edge $\{x, y\} \in E(G)$ for any two points $x, y \in N_i$, such that x, and y appear consecutively in a clockwise traversal of Γ_i. Moreover, for every $z \in N_i$, let $z' \in N_\Delta$ be such that $z \in \gamma(z')$. Let z'' be the point in the intersection of $\gamma(z')$ with Γ_{i-1}. If $z'' \in N_{i-1}$, then we add the unit-length edge $\{z, z''\}$. Otherwise, let w be the first point in N_{i-1} that we visit in a clockwise traversal of Γ_{i-1} starting from z''. We add the unit-length edge edge

$\{z, w\}$. This concludes the definition of the graph G. It is straightforward to check that G is a funnel with basepoint x^*.

We can now define an embedding $f : X \to V(G)$, by mapping every points $x \in X$ to its nearest neighbor in $V(G)$. It remains to verify that f has constant distortion. Observe that the set $V(G)$ contains a $2r$-net in D_Δ, and therefore for any $x \in X$, we have $d(x, f(x)) < 2r$. Since the minimum distance in X is at least 1, this implies that f is an injection, and for any $x, y \in X$, we have $d(x, y) = \Theta(d(f(x), f(y)))$. It therefore suffices to show that for any $x, y \in V(G)$, we have $d_G(x, y) = \Theta(d(x, y))$.

We first show that for any $x, y \in V(G)$, we have $d_G(x, y) = \Omega(d(x, y))$. To that end, it suffices to show that for any edge $\{x, y\} \in E(G)$, we have $d(x, y) = O(d_G(x, y)) = O(1)$.

We consider first case where there exists $i \in \{1, \ldots, \Delta\}$ such that $x, y \in N_i$, and x, y are consecutive in Γ_i. Let α be the arc of Γ_i between x, and y, that does not contain any other points in N_i. By the triangle inequality, there exists $z \in \alpha$, such that $d(x, z) \geq d(x, y) \geq 2$, and $d(y, z) \geq d(x, y) \geq 2$. Since N_i is an r-net in Γ_i, it follows that there exists $z' \in N_i$, such that $d(z, z') < r$. Let β be the geodesic between z, and z'. The arc β intersects either $\gamma(x)$, or $\gamma(y)$. Assume w.l.o.g. that it intersects $\gamma(x)$ at some points z''. By lemma 2.2 we have that as we travel along β, the distance to x^* is a convex function. This implies that $d(x, z'') \leq d(z, z')$. We conclude that $d(x, y) \leq 2d(x, z) \leq 2(d(x, z'') + d(z'', z)) \leq 2(d(x, z'') + d(z', z)) \leq 4d(z, z') \leq 4r = O(1)$.

Next, we consider the case where $x \in N_i$, and $y \in N_{i+1}$, for some $i \in \{0, \ldots, \Delta\}$. Let y' be the point where $\gamma(y)$ intersects Γ_i. Arguing as above, we have that $d(y', x) = O(1)$. Therefore, $d(x, y) \leq d(y, y') + d(y', x) \leq r + O(1) = O(1)$. This concludes that proof that for any edge $\{x, y\} \in E(G)$, we have $d(x, y) = O(1)$, and therefore for any $x, y \in V(G)$, we have $d_G(x, y) = \Omega(d(x, y))$.

It remains to show that for any $x, y \in V(G)$, we have $d_G(x, y) = O(d(x, y))$. We consider first the case where there exists $i \in \{1, \ldots, \Delta\}$, such that $x, y \in N_i$ (the case $i = 0$ is trivial since N_0 contains only x^*). Let β be a geodesic between x, and y. By lemma 2.2, we have $\beta \subset D_i$. Let x' be the unique point in $\gamma(x) \cap \Gamma_{i-1}$, and let y' be the unique point in $\gamma(y) \cap \Gamma_{i-1}$. By lemma 2.3 we have $d(x', y') \leq d(x, y)$. Let x'' be the parent of x, and let y'' be the parent of y in G. Let $x' = z_1, \ldots, z_k = y'$ be the points in N_{i-1} that appear between x', and y' along Γ_{i-1}. For any $i \in \{1, \ldots, k\}$, pick a child w_i of z_i, with $w_1 = x$, and $w_k = y$. For any $i \in \{1, \ldots, k\}$, the curve β intersects $\gamma(w_i)$. By the above discussion we have that the distance between any two such consecutive intersection points is $\Omega(1)$. Therefore, $d(x, y) = \mathsf{len}(\beta) = \Omega(k)$. The x-y path in G that visits the vertices $xz_1 \ldots z_k y$ in this order has length $k + 2$, and therefore $d_G(x, y) = O(d(x, y))$.

Next, we consider the case where there exists $i \in \{1, \ldots, \Delta\}$, such that $x \in N_i$, and $y \in N_{i-1}$. This case is identical to the case above, by replacing y with y''. We therefore also obtain $d_G(x, y) = O(d(x, y))$ in this case.

Finally, we consider the case of arbitrary points $x, y \in V(G)$. Let β be the geodesic between x, and y. The curve β can be decomposed into consecutive segments $\beta_1, \ldots, \beta_k$, such that every such segment is contained in (the closure of) $D_i \setminus D_{i-1}$, for some $i \in \{1, \ldots, \Delta\}$. Consider such a segment β_i. There exists $j, \ell \in \{0, \ldots, \Delta\}$, with $|j - \ell| \leq 1$, and such that $x_i \in \Gamma_j$, and $y_i \in \Gamma_\ell$. Let x'_i be the nearest neighbor of x_i in N_j, and let y'_i be the nearest neighbor of y_i in N_ℓ. Since N_j is a $O(1)$-net for Γ_j, and N_ℓ is a $O(1)$-net for Γ_ℓ, we have $d(x'_i, y'_i) \leq d(x_i, y_i) + O(1) = O(d(x_i, y_i))$. By the above analysis we have $d_G(x'_i, y'_i) = O(d(x'_i, y'_i))$. Therefore, we obtain $d_G(x_i, y_i) = O(d(x_i, y_i))$. We conclude that $d_G(x, y) \leq \sum_i d_G(x_i, y_i) = O(\sum_i d(x_i, y_i)) = O(d(x, y))$, as required. $\square$

3 Cutting along a ray

We now show that every funnel admits a constant-distortion stochastic embedding into a distribution over simpler graphs, that we call *pyramids*. Intuitively, a pyramid is obtained by "cutting" a funnel along a ray. The structure of pyramids will simplify the exposition of the embedding into L_1 that we describe in the subsequent sections.

Definition 3.1 (Pyramid)**.** *Let G be an unweighted planar graph, let $v \in V(G)$, and let $\Delta \geq 1$ be an integer. We say that G is a* pyramid *with* basepoint v, *and of* depth Δ *if the following conditions are satisfied:*

(1) There exists a collection of pairwise vertex-disjoint paths $P_1, \ldots, P_\Delta \subset G$, with $P_i = u_{i,1} \ldots u_{n_i,i}$, such that $V(G) = \bigcup_{i=1}^{\Delta} V(P_i)$. For notational convenience, we allow a path P_i to consist of a single vertex, in which case it has no edges. Moreover, we have $V(P_1) = \{v\}$. We refer to each P_i as a layer.

(2) For every $i \in \{2, \ldots, \Delta - 1\}$, the graph $G \setminus V(P_i)$ has exactly two connected components, one with vertex set $\bigcup_{j=1}^{i-1} V(P_j)$, and another with vertex set $\bigcup_{j=i+1}^{\Delta} V(P_j)$.

(3) For every $i \in \{2, \ldots, \Delta\}$, every $u \in V(P_i)$ has exactly one neighbor u in $V(P_{i-1})$. We refer to this neighbor as the parent *of u. In particular, v is the parent of all vertices in $V(P_2)$.*

(4) For every $i \in \{1, \ldots, \Delta - 1\}$, every $w \in V(P_i)$ at least one neighbor w' in $V(P_{i-1})$. We refer to every such w' as a child *of w.*

(5) For any $i \in \{1, \ldots, \Delta - 1\}$, and for any $\{u_{i,j}, u_{i+1,j'}\}, \{u_{i,t}, u_{i+1,t'}\} \in E(G)$, we have $j \leq t \iff j' \leq t'$. In other words, the ordering of the vertices in P_{i+1} agrees with the ordering of their parents in P_i.

We say that a path R in G between v, and a vertex $u \in V(P_\Delta)$, is a ray. *We denote by* $\mathsf{Pyramids}$ *the family of all pyramid graphs. Figure 2 depicts an example of a pyramid.*

Definition 3.2 (Skeleton of a pyramid)**.** *Let G be a pyramid with basepoint $v \in V(G)$. We define the* skeleton *of G to be a tree T, with $V(T) = V(G)$, with root v, and with*

$$E(T) = \left\{ \{x, y\} \in \binom{V(G)}{2} : x \text{ is the parent of } y \right\}.$$

For any $x, y \in V(G)$, we denote by nca *the nearest common ancestor of x, and y in T. We also define for any $x \in V(G)$,*

$$\mathsf{depth}(x) = d_T(v, x) + 1.$$

Figure 2 depicts an example of a skeleton.

Definition 3.3 ($\prec$)**.** *For any $i \in \{1, \ldots, \Delta\}$, for any $u_{i,j}, u_{i,j'} \in V(P_i)$, with $j < j'$, we write $u_{i,j} \prec u_{i,j'}$. Moreover, for any $x, y \in V(G)$, such that x, and y do not lie on the same ray, let $z = \mathsf{nca}(x, y)$, and let x' (resp. y') be the child of z in the z-x (resp. z-y) path in T. Then, we write $x \prec y$ if and only if $x' \prec y'$. Finally, for any $x'', y'' \in V(G)$, we write $x'' \preceq y''$ if and only if either $x'' \preceq y''$, or x'', and y'' lie on the same ray.*

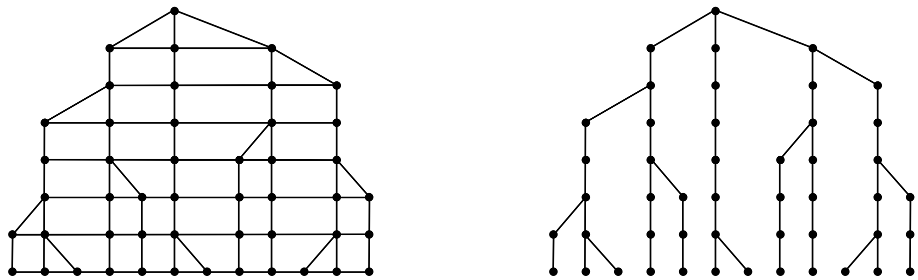

Figure 2: A pyramid (left), and its skeleton (right).

Lemma 3.4 (Pyramid representation)**.** *For every funnel G, we have $G \rightsquigarrow \oplus_1\{\mathsf{Pyramids}\}$.*

Proof. Let G be a funnel with basepoint $x^* \in V(G)$, and depth Δ. Let R be a ray in G. Replace $R \backslash x^*$ by a $\Delta \times 4$ grid H. Clearly, this results into an embedding of G into a funnel G' with distortion $O(1)$. Let R' be the union of the two central columns of H, and let $A = R' \cup \{x^*\}$. Observe that $\mathsf{dil}_{G'}(A) = 1$. Applying lemma 1.5 on G' and the set A, we obtain a stochastic embedding of G' into a distribution of graphs $\mathcal{D}$. Since G' is planar, and $\mathsf{dil}_{G'}(A) = 1$, it follows by Theorem 1.3 that the distortion of the resulting stochastic embedding is $O(1)$. Every graph in the support of $\mathcal{D}$ is obtained via 1-sums of $G'[A]$, with $G'[V \setminus A \cup \{a\}]$, for some $a \in A$. The graph $G'[A]$ is a $\Delta \times 2$ grid, with the basepoint x^* connected to the two vertices in the top row, and is therefore a pyramid. For any $a \in A$, the graph $G'[V \setminus A \cup \{a\}]$ is obtained from G' by cutting along a ray, and is therefore also a pyramid. This concludes the proof. $\square$

4 Monotone cuts

In this section we describe the family of cuts, that we will use when defining our embedding into L_1. These are cuts that we call *monotone*, and intuitively correspond to sets that only cross every ray at most once. We also describe a specific "shifting" operation that will allow us to modify a cut in order to adapt to the finer geometry of a given space.

Definition 4.1 (Monotone cut)**.** *Let G be a pyramid with basepoint $v \in V(G)$, and let $S \subseteq V(G)$. We say that S is v-*monotone *(or monotone when v is clear from the context) if $v \in S$, and for any ray R in G, $R \cap S$ is a prefix of R. In particular, this implies that $G[S]$ is a connected subgraph (see Figure 3).*

Definition 4.2 (Boundary of a monotone cut)**.** *Let $S \subseteq V(G)$ be a monotone cut. We define the* vertex boundary *of S, denoted by $\partial_V S$, to be the set of all $u \in S$, such that all children of u are not in S. We also define the* edge boundary *of S, denoted by $\partial_E S$, to be*

$$\partial_E S = \{\{x, y\} \in E(G) : x, y \in \partial_V S \text{ and } \mathsf{depth}(x) = \mathsf{depth}(y)\}.$$

Finally, we define the graph $\partial S = (\partial_V S, \partial_E S)$ (see Figure 3).

Definition 4.3. *Let G be a pyramid, let T be the skeleton of G. Let $u \in V(G)$, and $r \geq 0$. Then, we denote by $\widetilde{N}(u, r)$ the set of all vertices $w \in V(G)$, such that u is an ancestor of w in T, and $d_T(u, w) \leq r$.*

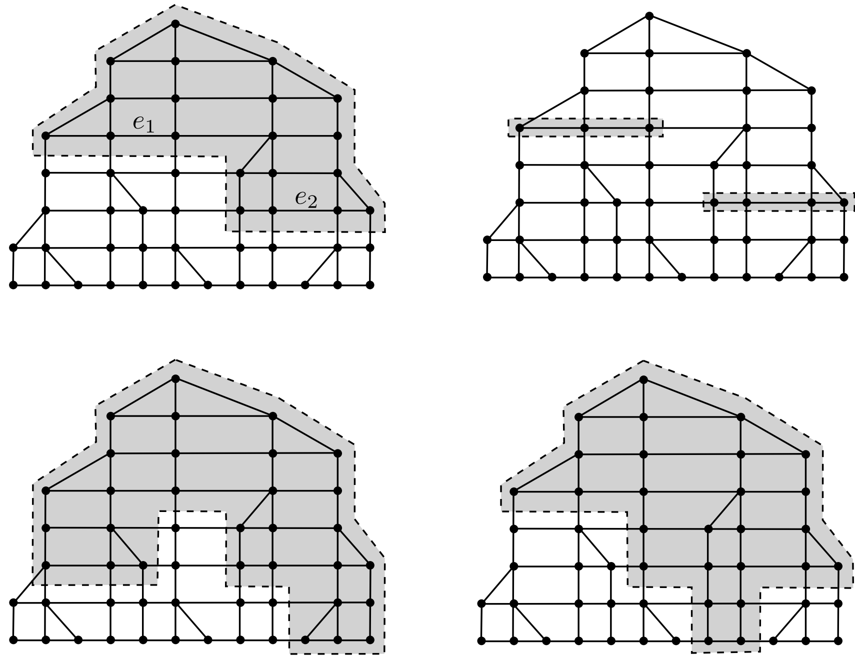

Figure 3: A monotone cut (top-left), its boundary (top-right), its odd $(2, \{e_1, e_2\})$-shift (bottom-left), and its even $(2, \{e_1, e_2\})$-shift (bottom-right).

Definition 4.4 (Odd/even shift of a monotone cut)**.** *Let $S \subseteq V(G)$ be a monotone cut, let $r > 0$, and $Z \subseteq \partial_E S$. Let $Z = \{\{x_i, y_i\}\}_{i=1}^k$, with*

$$x_1 \prec y_1 \preceq x_2 \prec y_2 \preceq \ldots \preceq x_k \prec y_k.$$

Let $\{V_i\}_{i=1}^{k+1}$ be a decomposition of $\partial_V S$, with $V_1 = \{u \in \partial_V S : u \preceq x_1\}$, $V_{k+1} = \{u \in \partial_V S : y_k \preceq u\}$, and for any $i \in \{2, \ldots, k\}$, $V_i = \{u \in \partial_V S : y_i \preceq u \preceq x_{i+1}\}$. We define a partition $\partial_V S = V_{odd} \cup V_{even}$, by setting $V_{odd} = \bigcup_{i=1}^{\lceil t/2 \rceil} V(Q_{2i-1})$, and $V_{even} = \bigcup_{i=1}^{\lfloor t/2 \rfloor} V(Q_{2i})$. We define the odd (r, Z)-shift *of S to be the cut S_{odd} given by*

$$S_{odd} = S \cup \bigcup_{u \in V_{odd}} \widetilde{N}(u, r).$$

Similarly, we define the even (r, Z)-shift *of S to be the cut S_{even} given by*

$$S_{even} = S \cup \bigcup_{u \in V_{even}} \widetilde{N}(u, r).$$

We say that a cut S' is a (r, Z)-shift of S, if it is either the odd, or the even (r, Z)-shift of S (see Figure 3 for an example).

5 The embedding

In this section we present a constant-distortion embedding of pyramids into L_1. Combining with lemmas 1.4, 2.4, & 3.4, this implies that every planar metric of non-positive curvature embeds into

L_1 with constant distortion.

Let G be a pyramid, with basepoint $v \in V(G)$. Let $\Delta \geq 1$ be the depth of G, and let $\delta = \lceil \log \Delta \rceil$. It will be convenient for our exposition to isometrically embed G into a larger pyramid G', with depth $\Delta' = O(\Delta)$, as follows. The pyramid G' contains a copy of G, and a new basepoint v', that is connected to the basepoint v of G via a path of length 2Δ, resulting into a pyramid of depth $\Delta' = 3\Delta$.

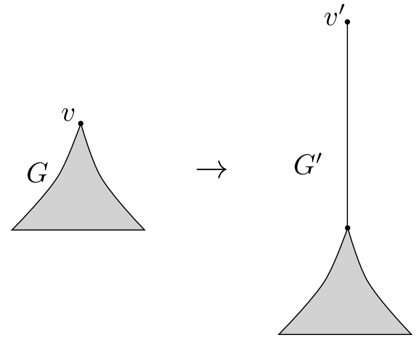

We will then compute an embedding for G', and prove that its restriction on G has small distortion. We remark that our embedding will have unbounded distortion for points in G' close to v' (more precisely, pairs of vertices at distance ε from v', will be distorted by a factor of $O(1/\varepsilon)$). However, this does not affect our result, since we only case about distances in G, which lies far from v'.

Definition 5.1 (Evolution of a monotone cut)**.** *Let $r > 0$, and let $S \subseteq V(G')$ be a monotone cut. The r-*evolution *of S is a probability distribution $\mathcal{D}$ over monotone cuts, defined by the following random process. Let*

$$Y = \{\{x, y\} \in \partial_E S : \mathsf{depth}(x) - \mathsf{depth}(\mathsf{nca}(x, y)) \in [r, 6r)\}.$$

Pick a random subset $Y' \subseteq Y$, by choosing every $e \in Y$ independently, with probability $1/r$. We probability $1/2$, let S' be the odd (r, Y')-shift of S, and otherwise let S' be the even (r, Y')-shift of S. The resulting random cut S' defines the distribution $\mathcal{D}$.

Let $\mathcal{M}$ be the set of all v-monotone cuts in G'. We inductively define a sequence $\{\mu_i\}_{i=0}^{\delta+1}$, where each μ_i is a probability distribution over $\mathcal{M}$. We define μ_0 as follows. Let $P_1, \ldots, P_{\Delta'}$ be the layers of G'. For any $j \in \{1, \ldots, \Delta'\}$, let $X_i = \bigcup_{t=1}^{j} V(P_t) = \mathsf{ball}(v', i)$. Let μ_0 be the uniform distribution over the collection of cuts $X_1, \ldots, X_{\Delta'}$.

For any $i \geq 0$, given μ_i, we inductively define μ_{i+1} via the following random process: We first pick a random cut S_i according to μ_i. Let $\mathcal{D} = \mathcal{D}(S_i)$ be the $\Delta/3^i$-evolution of S_i. We pick a random cut S_{i+1} according to $\mathcal{D}$. The resulting random variable S_{i+1} defines the probability distribution μ_{i+1}.

We define the embedding f induced by the probability distribution μ_δ, and the embedding f_0 induced by the probability distribution μ_0. Finally, we set the resulting embedding to be

$$g = f \oplus f_0,$$

i.e. the concatenation of the embeddings f, and f_0. In the next section we show that the distortion of g restricted on G is bounded by some universal constant.

6 Distortion analysis

We now analyze the distortion of the embedding g constructed in the previous section.

6.1 Distortion of vertical pairs of points

Lemma 6.1. *Let $u \in V(G)$, with $\mathsf{depth}(u) < \Delta'$, and let $i \in \{1, \ldots, \delta\}$. Then,* $\Pr[u \in \partial_V S_i] = 1/\Delta'$.

Proof. The proof is by induction on i. For $i = 0$, the assertion holds since μ_0 is the uniform distribution over the cuts $X_1, \ldots, X_{\Delta'}$. Suppose next that $i > 0$. Let $r = \Delta/3^{i-1}$, and let u' be the ancestor of u in T, with $d_T(u, u) = r$. Fix some S_{i-1} in the support of μ_{i-1}, and suppose that S_i is sampled from the r-evolution of S_{i-1}. This means that we first sample a set of edges Y, and for any such Y we set S_i to be the odd (r, Y)-shift of S_{i-1} with probability $1/2$, or otherwise we set S_i to be the even (r, Y)-shift of S_{i-1}. Therefore, we have have $u \in \partial_V S_i$, only if either $u \in \partial_V S_{i-1}$, or $u' \in \partial_V S_{i-1}$. Conditioned on either of these two events, and for any Y, exactly one of the odd/even shifts of S_{i-1} has u in its boundary. This implies that

$$\begin{aligned}
\Pr[u \in \partial_V S_i] &= \Pr[u \in \partial_V S_i | u \in \partial_V S_{i-1}] \cdot \Pr[u \in \partial_V S_{i-1}] \\
&\quad + \Pr[u \in \partial_V S_i | u' \in \partial_V S_{i-1}] \cdot \Pr[u' \in \partial_V S_{i-1}] \\
&= \frac{1}{\Delta'} \cdot \frac{1}{2} + \frac{1}{\Delta'} \cdot \frac{1}{2} \\
&= 1/\Delta',
\end{aligned}$$

as required. $\square$

Lemma 6.2. *Let $x, y \in V(G)$, such that x, y lie on the same ray. Then,* $\|f(x) - f(y)\|_1 = d_G(x, y)/\Delta$.

Proof. Let R be the ray containing both x, and y. Let R' be the subpath of R between x, and y, including x, and excluding y. By the monotonicity of $\mathcal{S}_\delta$, it follows that $\mathbf{1}_{S_\delta}(x) \neq \mathbf{1}_{S_\delta}(y)$, if and only if there exists $z \in V(R')$, such that $z \in \partial_V S_\delta$. Since these events are disjoint for different z, we obtain by lemma 6.1 that $\|f(x) - f(y)\|_1 = \Pr[\mathbf{1}_{S_\delta}(x) \neq \mathbf{1}_{S_\delta}(y)] = |V(R')|/\Delta' = d_G(x, y)/\Delta'$, as required. $\square$

6.2 Distortion of horizontal pairs of points

We now bound the distortion on pairs of vertices $x, y \in V(G)$ that lie on the same layer of G', i.e. such that $\mathsf{depth}(x) = \mathsf{depth}(y) = h$. Let $d_G(x, y) = L$. Let also $h' = \mathsf{depth}(\mathsf{nca}(x, y))$. We assume w.l.o.g. that $x \preceq y$. Let P be the subpath of P_h between x, and y.

Let

$$E_{\text{top}} = \{\{z, w\} \in E(P) : \mathsf{depth}(\mathsf{nca}(z, w)) \leq h - L/2\},$$

and

$$E_{\text{bottom}} = E(P) \setminus E_{\text{top}}.$$

Lemma 6.3. $|E_{top}| \leq L$.

Proof. Suppose, to the contrary, that $|E_{\mathrm{top}}| > L$. For any $i \in \{h - L/2, \ldots, h\}$, let Z_i be the subpath of P_i between the ancestor of x, and the ancestor of y in P_i. For any $e = \{z, w\} \in E_{\mathrm{top}}$, with $z \prec w$, let R_z be a ray containing z, and let W_e be the subpath of R_e contained between $P_{h-L/2}$, and P_h.

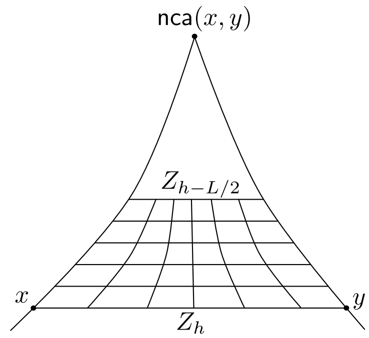

The union of all these paths $(\bigcup_i Z_i) \cup (\bigcup_e W_e)$ forms a $(L/2+1) \times L$ grid minor in G', with x, and y being the bottom-left, and bottom-right vertices respectively. Since the x-y shortest path in G' is contained in $\mathsf{ball}(v', h)$, this implies that $d_G(x, y) > L$, which is a contradiction. $\square$

Let H be the subgraph of G induced on the set of vertices

$$V(H) = \{u \in V(G) : h' \leq \mathsf{depth}(u) \leq h \text{ and } x \preceq u \preceq y\}.$$

Definition 6.4 (Straight cut)**.** *Let $i \in \{1, \ldots, \delta\}$, and $j \in \{1, \ldots, \Delta'\}$. We say that S_i is j-*straight *if $\partial S_i \cap H \subseteq P_j$.*

Let $e = \{z, w\} \in E(P)$. We say that an edge $e' = \{z', w'\} \in E(G)$ is an *ancestor* of e, if z' is an ancestor of z in T, w' is an ancestor of w in T, and $\mathsf{depth}(z') = \mathsf{depth}(w')$.

Definition 6.5 (Bend)**.** *Let $e \in E(P)$. We say that e* bends *S_i, if the following events happen.*

(1) There exists $j \in \{1, \ldots, \Delta'\}$, such that S_i is j-straight.

(2) Let $Y \subseteq \partial_E S_i$, such that S_{i+1} is the (r, Y)-shift of S_i, for some $r > 0$. Then, there exists an ancestor of e in Y.

Lemma 6.6. *Let $j \in \{h', \ldots, h\}$, and let $i \in \{1, \ldots, \delta\}$. Then,* $\Pr[S_i$ *is j-straight*$] \leq 1/\Delta'$.

Proof. Let z be an arbitrary vertex in $V(P_j) \cap V(H)$. Clearly, S_i can only be j-straight if $z \in \partial_V S_i$. Therefore, by lemma 6.1 we obtain $\Pr[S_i$ is j-straight$] \leq \Pr[z \in \partial_V S_i] = 1/\Delta'$, as required. $\square$

For any edge $e = \{z, w\} \in E(P)$, and for any $i \in \{1, \ldots, \delta\}$, let $\mathcal{E}(e, i)$ be the conjunction of the following two events:

- $\mathcal{E}_1(e, i)$: There exists j, such that the following event, denoted by $\mathcal{E}_1(e, i, j)$, holds: Intuitively, the event $\mathcal{E}_1(e, i, j)$ describes a necessary condition such that a bend of S_i can potentially lead to a cut S_δ that separates x, and y. Formally, we have that S_i is j-straight, with

$$\Delta/3^i \leq j - \mathsf{depth}(\mathsf{nca}(z, w)) < 6\Delta/3^i, \tag{4}$$

and

$$h - j \leq 2\Delta/3^i. \tag{5}$$

- $\mathcal{E}_2(e,i)$: e bends S_i.

Lemma 6.7. *Suppose that $\mathbf{1}_{S_\delta}(x) \neq \mathbf{1}_{S_\delta}(y)$. Then, there exists $e \in E(P)$, and $i \in \{1,\ldots,\delta\}$, such that the event $\mathcal{E}(e,i)$ occurs.*

Proof. Recall that by the definition of μ_0, the cut S_0 is j_0-straight, for some $j_0 \in \{1,\ldots,\Delta'\}$. Since $\mathbf{1}_{S_\delta}(x) \neq \mathbf{1}_{S_\delta}(y)$, it follows that for all $j_\delta \in \{1,\ldots,\Delta'\}$, the cut S_δ is not j_δ-straight. Let $i^* \in \{0,\ldots,\delta-1\}$ be the smallest integer such that for all $j \in \{1,\ldots,\Delta'\}$, the cut S_{i^*+1} is not j-straight. This means that S_{i^*} is j^*-straight, for some $j^* \in \{1,\ldots,\Delta'\}$. Therefore, there exists $e=\{z,w\} \in E(P)$, such that e bends S_{i^*}, which means that the event $\mathcal{E}_2(e,i^*)$ occurs. It suffices to show that the event $\mathcal{E}_1(e,i^*,j^*)$ also occurs. We have established that S_{i^*} is j^*-straight, so its remains to show that (4) & (5) hold. Condition (4) follows immediately form the fact that e bends S_{i^*}, and S_{i^*+1} is the $(Y,\Delta/3^{i^*})$-shift of S_{i^*}, with $e \in Y$. Since S_{i^*} is j^*-straight, we have $S_{i^*} \subseteq \mathsf{ball}(v',j^*)$. The cut S_δ is obtained from S_{i^*} via a sequence of (Y,r)-shifts, with exponentially decreasing values of r. This implies $S_\delta \subseteq \mathsf{ball}(v',t)$, for some $t \leq j^* + \sum_{i=i^*}^{\delta} \Delta/3^i < j^* + 2\Delta/3^*$. Since $\mathbf{1}_{S_\delta}(x) \neq \mathbf{1}_{S_\delta}(y)$, we have $t > h$, and therefore $h - j^* < 2\Delta/3^{i^*}$, which implies (5), and concludes the proof. □

Lemma 6.8 (Expansion of horizontal pairs)**.** *Let $x,y \in V(G)$, such that $\mathsf{depth}(x) = \mathsf{depth}(y)$. Then, $\|f(x)-f(y)\|_1 = O(d(x,y)/\Delta')$.*

Proof. Let $\mathcal{E}_{\mathrm{top}}$ denote the event that there exists $e \in E_{\mathrm{top}}$, and $i \in \{1,\ldots,\delta\}$, such that $\mathcal{E}(e,i)$ occurs. Similarly, let $\mathcal{E}_{\mathrm{bottom}}$ denote the event that there exists $e \in E_{\mathrm{bottom}}$, and $i \in \{1,\ldots,\delta\}$, such that $\mathcal{E}(e,i)$ occurs. By lemma 6.7 we have

$$\begin{aligned}\|f(x)-f(y)\|_1 &= \Pr[\mathbf{1}_{S_\delta}(x) \neq \mathbf{1}_{S_\delta}(x)] \\ &\leq \Pr[\mathcal{E}_{\mathrm{top}}] + \Pr[\mathcal{E}_{\mathrm{bottom}}].\end{aligned}$$

Let us bound the two latter quantities separately.

We first bound $\Pr[\mathcal{E}_{\mathrm{top}}]$. Let $e=\{z,w\} \in E_{\mathrm{top}}$, and $i \in \{1,\ldots,\delta\}$. Let $h' = \mathsf{depth}(\mathsf{nca}(\{z,w\}))$. Recall that by the definition of $\mathcal{E}_1(e,i,j)$, in order for $\mathcal{E}_1(e,i,j)$ to occur for some j, we must have by (4) that $j - h' \leq 6\Delta/3^i$, and by (5) that $h - j \leq 2\Delta/3^i$. We therefore obtain that

$$h - h' = h - j + j - h' = O(\Delta/3^i).$$

Note that S_{i+1} is the (Y,r)-shift of S_i, for some $r = \Delta/3^i$, and for random some $Y \subseteq \partial_E S_i$. In order for the edge e to bend S_i, its must be the case that its unique ancestor (if it exists) in $\partial_E S_i$ is chosen in Y. Every edge in chosen in Y with probability at most $1/r$. Therefore, for any j, and for any i, we have

$$\Pr[\mathcal{E}_2(e,i)|\mathcal{E}_1(e,i,j)] \leq 3^i/\Delta.$$

Moreover, $\mathcal{E}_1(e,i,j)$ can occur only if $j \in \{h',\ldots,h\}$. For each such value $j \in \{h',\ldots,h\}$, and for any i, we have by lemma 6.6 that

$$\Pr[\mathcal{E}_1(e,i,j)] = O(1/\Delta).$$

To summarize, we have

$$
\begin{aligned}
\Pr[\mathcal{E}_{\text{top}}] &\leq \sum_{e \in E_{\text{top}}} \sum_{i \in \{1,\ldots,\delta\}} \Pr[\mathcal{E}(e,i)] \\
&\leq \sum_{e \in E_{\text{top}}} \sum_{j \in \{h',\ldots,h\}} \sum_{i \in \{1,\ldots,\delta\}} \Pr[\mathcal{E}_2(e,i)|\mathcal{E}_1(e,i,j)] \cdot \Pr[\mathcal{E}_1(e,i,j)] \\
&\leq \sum_{e \in E_{\text{top}}} \sum_{j \in \{h',\ldots,h\}} \sum_{i \in \{1,\ldots,\delta\}} \frac{3^i}{\Delta} O(1/\Delta) \\
&\leq \sum_{e \in E_{\text{top}}} \sum_{j \in \{h',\ldots,h\}} O(1/(h-h')) \cdot O(1/\Delta) \\
&\leq \sum_{e \in E_{\text{top}}} O(1/\Delta) \\
&= O(|E_{\text{top}}|/\Delta) \\
&= O(L/\Delta').
\end{aligned} \tag{6}
$$

We next bound $\Pr[\mathcal{E}_{\text{bottom}}]$. Let $\{1,\ldots,\delta\}$, $j \in \{1,\ldots,\Delta'\}$, $e \in E_{\text{bottom}}$, such that both $\mathcal{E}_1(e,i,j)$, and $\mathcal{E}_2(e,i)$ occur. As above, let $e = \{z,w\}$, and $h' = \mathsf{depth}(\mathsf{nca}(z,w))$. Then, we must have $h' \leq j \leq h$, which implies $j - h' \leq h - h' \leq L/2$. Since S_{i+1} is the (r,Y)-shift of S_i for some $Y \subseteq E(P)$, with $r = \Delta/3^i$, we obtain that $j - h' \in [r, 6r)$, which implies $3^i \geq 2\Delta/L$. Let R_x be the ray containing x, and let χ be the unique vertex in the intersection of R_x with ∂S_i. Let also χ' be the unique vertex in the intersection of R_x with ∂S_δ. For every $i' \geq i$, the intersection of $\partial S_{i'}$ with R_x moves by at most $\Delta/3^{i'}$ along R_x, and therefore $d_T(\chi,\chi') < 2\Delta/3^{i'} = O(L)$. Since $\chi \in P_j$, and $j \in [h', h]$, it follows that $\mathsf{depth}(\chi)$ can take at most $h' - h + 1$ different values. Therefore, χ' can only lie inside a subpath $R'_x \subseteq R_x$ of length $O(h'-h)$. Applying lemma 6.1, we obtain

$$
\begin{aligned}
\Pr[\mathcal{E}_{\text{bottom}}] &\leq \Pr[\chi' \in R'_x] \\
&\leq |V(R'_x)|/\Delta' \\
&= O(h'-h)/\Delta' \\
&= O(L/\Delta').
\end{aligned} \tag{7}
$$

Combining (6) & (7) we conclude that $\|f(x) - f(y)\|_1 = O(L/\Delta') = O(d(x,y)/\Delta')$, as required. $\square$

We now bound the contraction of f. For any $i \in \{1,\ldots,\delta\}$, let

$$
J_i = \left\{ h - \frac{\Delta}{3^i}, \ldots, h - \frac{2}{3} \cdot \frac{\Delta}{3^i} \right\}.
$$

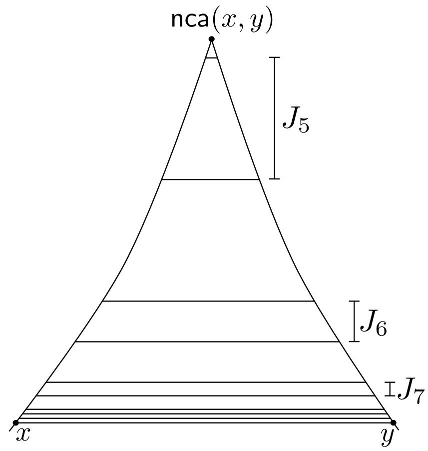

Lemma 6.9. *Let $i \in \{1, \ldots, \delta\}$, $j \in J_i$, such that S_i is j-straight. Let $Y \subseteq E(G')$, such that S_{i+1} is the $(\Delta/3^i, Y)$-shift of S_i, where $|Y \cap E(P)|$ is odd. Then, $\mathbf{1}_{S_\delta}(x) \neq \mathbf{1}_{S_\delta}(y)$.*

Proof. Since $j \in J_i$, it follows that $j \leq h - \Delta/3^i$. Let x' be the ancestor of x with $d_G(x, x') = \Delta/3^i$, and let y' be the ancestor of y with $d_G(y, y') = \Delta/3^i$. Since S_i is j-straight, it follows that $x', y' \in S_i$. Suppose that S_{i+1} is an odd shift of S_i; the case where S_{i+1} is an even shift of S_i is completely symmetric by exchanging x, and y. Since S_{i+1} is the odd $(\Delta/3^i, Y)$-shift of S_i, and $|Y \cap E(P)|$ is odd, it follows that $x \in S_{i+1}$, and $y \notin S_{i+1}$. Since $S_\delta \supset \ldots \supset S_{i+1}$, we obtain $x \in S_\delta$. Let W be the subpath of the ray containing y, and y'. We have that for any $k > 1$, the cut S_{i+k} contains a prefix of W of length at most $\Delta/3^{i+1} + \ldots + \Delta/3^\delta < 2\Delta/3^{i+1} \leq h - j = d_G(y, y')$, and therefore y is not in S_δ, concluding the proof. □

For any $t \in \{h', \ldots, h\}$, let

$$E_t = \{\{z, w\} \in E(P) : \mathsf{depth}(\mathsf{nca}(z, w)) \leq t\}.$$

Lemma 6.10. *There exists $t^* \in \{h - L/2, \ldots, h - L/4\}$, such that $L/2 \leq |E_{t^*}| \leq L$.*

Proof. By lemma 6.3 we have $|E_{h-L/2}| \leq L$. Let x' be the ancestor of x with $d_G(x, x') = L/4$, and let y' be the ancestor of y with $d_G(y, y') = L/4$. We have $|E_{h-L/4}| \geq d_G(x', y') \geq d_G(x, y) - d_G(x, x') - d_G(y, y') \geq L/2$, and the assertion follows. □

For any $i \in \{1, \ldots, \delta\}$, and for any $j \in \{1, \ldots, \Delta'\}$, let

$$B_{i,j} = \bigcup_{t=j-6\Delta/3^i}^{j-\Delta/3^i} E_t$$

Intuitively, the set $B_{i,j}$ contains all edges in H that could possibly bend S_i, when S_i is j-straight.

Lemma 6.11. *There exists $i^* \in \{1, \ldots, \delta\}$, with $\Delta/3^{i^*} = \Omega(L)$, and such that*

$$\bigcup_{i=1}^{i^*-1} \bigcup_{j \in J_i} B_{i,j} = O(L), \tag{8}$$

and

$$\bigcup_{i=1}^{i^*} \bigcup_{j \in J_i} B_{i,j} = \Omega(L). \tag{9}$$

Proof. It is straightforward to verify that

$$\bigcup_{i=1}^{\delta} \bigcup_{j \in J_i} B_{i,j} = E(P).$$

Let t^* be as in lemma 6.10. It follows that by setting

$$i^* = \min\{i \in \{1, \ldots, \delta\} : h - \Delta/3^i \leq t^*\},$$

we have $E_{t^*} \subseteq \bigcup_{i=1}^{i^*} \bigcup_{j \in J_i} B_{i,j}$, and therefore conditions (8) & (9) are satisfied. Moreover, we have $\Delta/3^{i^*} \geq h - t^* \geq L/4 = \Omega(L)$, as required. $\square$

Lemma 6.12. *Let $i \in \{1, \ldots, i^*\}$, and $j \in J_i$. Then,* $\Pr[S_i$ *is j-straight*$] = \Omega(1/\Delta')$.

Proof. We use a coupling argument. Let $E^* = \bigcup_{i=1}^{i^*-1} \bigcup_{j \in J_i} B_{i,j}$. Our embedding uses a random process σ of sampling $S_1, \ldots, S_i$. We define a modified random process σ' for sampling monotone cuts $S'_1, \ldots, S'_i$ as follows. The process σ' uses the same algorithm as σ, with the only difference that when taking $S'_{\iota+1}$ to be the (r, Y)-shift of S'_ι, where Y is chosen from a set $E' = E'(S'_\iota)$ independently with probability $1/r$, we instead chose Y' from the set $E' \setminus E^*$. In other words, we execute the same algorithm, but we always ignore the edges in E^* when computing shifts.

Since we ignore the edges in E^*, the final cut S'_i is always either j'-straight, for some $j' \leq j$, or it contains P_j. Arguing as in the proof of lemma 6.1, it is straightforward to verify that every vertex $u \in V(G)$ appears in $\partial_V S_i$ with probability $1/\Delta'$. Therefore, $\Pr[S'_i$ is j-straight$] = 1/\Delta'$.

We can now re-define the original random process σ as follows. At every step, when computing a (r, Y)-shift of some cut S, we first pick Y' according to the choices made in σ', and then we augment Y' to a new set Y by adding independently, and with probability $1/r$ every edge from E^* that was ignored in σ' at this step. It is straightforward to check that this definition results in the same process σ.

Let us say that the coupling of these two processes *fails* if at some step the processes σ, and σ' deviate. Recall that we obtain a set $S_{\iota+1}$ by taking a (r, Y)-shift of a set S_ι, for some $r = \Delta/3^\iota$. Every edge in E^* is eligible for appearing in such a set Y at most $O(1)$ times during the process. Moreover, since $\iota \leq i^*$, we have that every eligible edge is chosen in a set Y with probability $1/r = O(3^{i^*}/\Delta) = O(1/L)$. It follows that for any execution of σ', the coupling does not fail with at least some constant probability $q > 0$. Thus, $\Pr[S_i$ is j-straight$] \geq q \cdot \Pr[S'_i$ is j-straight$] = \Omega(1/\Delta')$, as required. $\square$

We will use the following simple fact about the parity of the sum of independent Bernoulli random variables.

Proposition 6.13. *There exists $c > 0$, such that the following holds. Let $p, k > 0$, and let $X_1, \ldots, X_k$ be a collection of independent Bernoulli random variables, such that for any $i \in \{1, \ldots, k\}$, we have* $\Pr[X_i = 1] = p$. *Then,* $\Pr\left[\sum_{i=1}^{k} X_i \text{ is odd}\right] > \min\{1/4, cpk\}$.

Lemma 6.14 (Contraction of horizontal pairs)**.** *Let $x, y \in V(G)$, such that* $\mathsf{depth}(x) = \mathsf{depth}(y)$*. Then,* $\|f(x) - f(y)\|_1 = \Omega(d(x,y)/\Delta')$*.*

Proof. For any $i \in \{1, \dots, i^*\}$, and for any $j \in J_i$, let $\mathcal{W}_1(i,j)$ denote the event that S_i is j-straight. By lemma 6.12 we have

$$\Pr[\mathcal{W}_1(i,j)] = \Omega(1/\Delta').$$

Let $\mathcal{W}_2(i,j)$ denote the event that there exists $Y \subseteq E(G')$, such that $|Y \cap E(P)|$ is odd, and S_{i+1} is the $(\Delta/3^i, Y)$-shift of S_i. Conditioned on the event that S_i is j-straight, we have that S_{i+1} is the $(\Delta/3^i, Y)$-shift of S_i, for some random $Y \subset E(G')$, with $Y \cap E(P) \subseteq B_{i,j}$, where every element of $B_{i,j}$ is chosen independently with probability $p = 3^i/\Delta$. Applying Proposition 6.13 we deduce that

$$\Pr[\mathcal{W}_2(i,j)|\mathcal{W}_1(i,j)] = \Omega(\min\{1/4, |B_{i,j}|3^i/\Delta\}) = \Omega(\min\{1/4, |E_j|/(h-j)\}).$$

Consider some $e = \{z, w\} \in B_{i,j}$, with $\mathsf{depth}(\mathsf{nca}(z,w)) = h''$. The edge e appears in $B_{i',j'}$, for at least $\Omega(h - h'')$ different values of $j' \in \bigcup_{i'=1}^{i^*} J_{i'}$. Arguing as in the proof of lemma 6.8, we can show that for every such value j', we have $h - j' = \Theta(h - h'')$. Therefore,

$$\sum_{i \in \{1,\dots,i^*\}} \sum_{j \in J_i} |B_{i,j}|/(h-j) = \Omega\left(\left|\bigcup_{i \in \{1,\dots,i^*\}} \bigcup_{j \in J_i} B_{i,j}\right|\right) = \Omega(L)$$

Combining the above with lemma 6.9, we obtain

$$\begin{aligned}
\|f(x) - f(y)\|_1 &= \Pr[\mathbf{1}_{S_\delta}(x) \neq \mathbf{1}_{S_\delta}(x)] \\
&\geq \sum_{i \in \{1,\dots,i^*\}} \sum_{j \in J_i} \Pr[\mathcal{W}_2(i,j)|\mathcal{W}_1(i,j)] \cdot \Pr[\mathcal{W}_1(i,j)] \\
&= \Omega(1/\Delta') \min\left\{1/4, \sum_{i \in \{1,\dots,\delta\}} \sum_{j \in J_i} |B_{i,j}|/(h-j)\right\} \\
&= \Omega(L/\Delta') \\
&= \Omega(d_G(x,y)/\Delta'),
\end{aligned}$$

as required. □

6.3 Distortion of general pairs of points

Lemma 6.15 (Embedding pyramids into L_1)**.** *There exists a universal constant $c > 1$, such that every pyramid graph admits an embedding into L_1 with distortion at most c.*

Proof. We will show that the embedding $g = f \oplus f_0$ has constant distortion on G. Let $x, y \in V(G)$. Assume w.l.o.g. that $\mathsf{depth}(x) \geq \mathsf{depth}(y)$. Let R_x be the ray containing x, and let x' be the unique vertex in R_x, with $\mathsf{depth}(x') = \mathsf{depth}(y)$. By lemmas 6.8 & 6.14 we have that there exist universal constants $\alpha > \beta > 0$, such that for any

$$\beta d_G(x', y)/\Delta \leq \|f(x') - f(y)\|_1 \leq \alpha d_G(x', y)/\Delta \tag{10}$$

Note that

$$d_G(x,x') = \mathsf{depth}(x) - \mathsf{depth}(x') = \mathsf{depth}(x) - \mathsf{depth}(y) \geq d_G(x,y). \tag{11}$$

Thus, we have

$$\begin{aligned}
\|f(x)-f(y)\|_1 &\leq \|f(x)-f(x')\|_1 + \|f(x')-f(y)\|_1 && (12)\\
&\leq d_G(x,y)/\Delta + \alpha d_G(x',y)/\Delta && (13)\\
&\leq d_G(x,y)/\Delta + \alpha d_G(x',x)/\Delta + \alpha d_G(x,y)/\Delta \\
&= (\alpha+1) d_G(x,x')/\Delta + \alpha d_G(x,y)/\Delta \\
&\leq (\alpha+1) d_G(x,y)/\Delta + \alpha d_G(x,y)/\Delta && (14)\\
&= (2\alpha+1) d_G(x,y)/\Delta, && (15)
\end{aligned}$$

where (12) follows by the triangle inequality, (13) by lemma 6.2 & (10), and (14) by (11). By (15) we have

$$\begin{aligned}
\|g(x)-g(y)\|_1 &= \|f(x)-f(y)\|_1 + \|f_0(x)-f_0(y)\|_1 \\
&\leq (2\alpha+1) d_G(x,y)/\Delta + d_G(x,y)/\Delta \\
&= (2\alpha+2) d_G(x,y)/\Delta. && (16)
\end{aligned}$$

This bounds the expansion of g. It remains to bound the contraction of g.

Let $\gamma = \frac{\beta}{4(2\alpha+1)}$. Assume first that $d_G(x,y) \geq \gamma d_G(x,y)$. We have

$$\begin{aligned}
\|g(x)-g(y)\|_1 &\geq \|f_0(x)-f_0(y)\|_1 \\
&= d_G(x',y)/\Delta \\
&\geq \gamma d_G(x,y)/\Delta && (17)
\end{aligned}$$

Next, assume that $d_G(x,y) < \gamma d_G(x,y)$. We have

$$\begin{aligned}
\|g(x)-g(y)\|_1 &\geq \|f(x)-f(y)\|_1 \\
&\geq \|f(x')-f(y)\|_1 - \|f(x)-f(x')\|_1 \\
&\geq \beta d_G(x',y)/\Delta - (2\alpha+1) d_G(x,x')/\Delta && (18)\\
&> (1-\gamma)\beta d_G(x,y)/\Delta - \gamma(2\alpha+1) d_G(x,y)/\Delta \\
&> \frac{1}{2}\beta d_G(x,y)/\Delta && (19)
\end{aligned}$$

where (18) follows by (10) & (15). Combining (17) & (19), we obtain that for all $x, y \in V(G)$

$$\|g(x)-g(y)\|_1 \geq \frac{\beta}{4(2\alpha+1)} d_G(x,y)/\Delta. \tag{20}$$

From (16) & (20) we conclude that the distortion of g is at most $4(2\alpha+1)(2\alpha+2)/\beta = O(1)$, concluding the proof. □

6.4 Proof of the main result

Combining the above results, we can now prove our main theorem.

Proof. Proof of theorem 1.2 Let (X, d) be a planar metric of non-positive curvature. Using lemma 2.4, the metric (X, d) admits an embedding into some funnel G with distortion $c_1 = O(1)$. Using lemma 3.4 we can find a stochastic embedding of G into a distribution $\mathcal{F}$ over pyramids with distortion $c_2 = O(1)$. By lemma 6.15 every pyramid in the support of $\mathcal{F}$ admits an embedding into L_1 with distortion $c_3 = O(1)$. Combining with lemma 1.4 we obtain that G admits an embedding into L_1 with distortion $c_2 c_3$. Therefore (X, d) admits an embedding into L_1 with distortion $\gamma = c_1 c_2 c_3 = O(1)$, concluding the proof. □

Acknowledgements The author thanks James R. Lee for sharing numerous insights on the planar embedding conjecture, and for explaining some of the previous work on embeddings of hyperbolic spaces.

References

[AR98] Yonatan Aumann and Yuval Rabani. An $O(\log k)$ approximate min-cut max-flow theorem and approximation algorithm. *SIAM J. Comput.*, 27(1):291–301 (electronic), 1998.

[BS] Mario Bonk and Oded Schramm. Embeddings of gromov hyperbolic spaces. *Geom. Funct. Anal*, 10:266–306.

[BS05a] Sergei Buyalo and Viktor Schroeder. A product of trees as universal space for hyperbolic groups. *arXiv.org*, September 2005.

[BS05b] Sergei Buyalo and Viktor Schroeder. Embedding of hyperbolic spaces in the product of trees. *Geometriae Dedicata*, 113(1):75–93, 2005.

[CGN+06] Chandra Chekuri, Anupam Gupta, Ilan Newman, Yuri Rabinovich, and Alistair Sinclair. Embedding k-outerplanar graphs into l_1. *SIAM J. Discrete Math.*, 20(1):119–136 (electronic), 2006.

[CJLV08] Amit Chakrabarti, Alexander Jaffe, James R. Lee, and Justin Vincent. Embeddings of topological graphs: Lossy invariants, linearization, and 2-sums. In *FOCS*, pages 761–770, 2008.

[CSW13] Chandra Chekuri, F. Bruce Shepherd, and Christophe Weibel. Flow-cut gaps for integer and fractional multiflows. *J. Comb. Theory, Ser. B*, 103(2):248–273, 2013.

[DL97] Michel Marie Deza and Monique Laurent. *Geometry of cuts and metrics*, volume 15 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, 1997.

[GNRS04] Anupam Gupta, Ilan Newman, Yuri Rabinovich, and Alistair Sinclair. Cuts, trees and l_1-embeddings of graphs. *Combinatorica*, 24(2):233–269, 2004.

[KL06] Robert Krauthgamer and James R. Lee. Algorithms on negatively curved spaces. In *FOCS*, pages 119–132, 2006.

[KPR93] Philip N. Klein, Serge A. Plotkin, and Satish Rao. Excluded minors, network decomposition, and multicommodity flow. In *STOC*, pages 682–690, 1993.

[LLR95] N. Linial, E. London, and Y. Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15(2):215–245, 1995.

[LP13] James R. Lee and Daniel Poore. On th e 2-sum embedding conjecture. In *SoCG*, 2013.

[LR99] Tom Leighton and Satish Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *J. ACM*, 46(6):787–832, 1999.

[LR10] James R. Lee and Prasad Raghavendra. Coarse differentiation and multi-flows in planar graphs. *Discrete & Computational Geometry*, 43(2):346–362, 2010.

[LS09] James R. Lee and Anastasios Sidiropoulos. On the geometry of graphs with a forbidden minor. In *STOC*, pages 245–254, 2009.

[OS81] Haruko Okamura and P.D. Seymour. Multicommodity flows in planar graphs. *Journal of Combinatorial Theory, Series B*, 31(1):75 – 81, 1981.

[Pap05] A. Papadopoulos. *Metric spaces, convexity and nonpositive curvature*. Irma Lectures in Mathematics and Theoretical Physics Series. European Mathematical Society Publishing House, 2005.

[Rao99] Satish Rao. Small distortion and volume preserving embeddings for planar and euclidean metrics. In *Symposium on Computational Geometry*, pages 300–306, 1999.

[TL97] W.P. Thurston and S. Levy. *Three-dimensional geometry and topology. 1 (1997)*. Princeton mathematical series. Princeton University Press, 1997.